

Mark 6 Next-Generation VLBI Data System

Alan Whitney, David Lapsley

MIT Haystack Observatory

Contact author: Alan Whitney, e-mail: awhitney@haystack.mit.edu

Abstract

The Mark 6 VLBI data system is being developed by MIT Haystack Observatory as a next-generation disk-based VLBI data system capable of supporting the goals of VLBI2010, with a maximum sustained recording rate of 16 Gbps writing to an array of 32 magnetic disks. The Mark 6 is based on COTS hardware and open-source code and is being designed to transition easily from the widely used Mark 5 system. A successful 16 Gbps per station VLBI demonstration experiment was conducted with Mark 6 in late 2011 as a proof-of-concept. Haystack Observatory is collaborating with the NASA/GSFC High-End Network Computing Group in the selection of high-performance COTS hardware platforms and with Conduant Corporation in the development of a high-performance disk module for Mark 6. Existing Mark 5 systems will be upgradable to Mark 6, and existing Mark 5 SATA modules will be upgradeable for compatibility with Mark 6. The Mark 6 system is projected to be available to the VLBI community in late 2012.

1. Introduction

The demand for increasing data rates for VLBI observations is particularly acute in two disciplines: 1) geodetic-VLBI observations, which must gather as much data as possible over a multi-GHz bandwidth in a period of 5-15 seconds and 2) mm-VLBI observations, which are typically starved for sensitivity, as well as having the need to gather as much data as possible during the atmospheric coherence period of 30-60 seconds. Both of these disciplines are eager to capture data at 16-64 Gbps and record to disk. The Mark 5 series of VLBI data systems is limited to a maximum of 4 Gbps (Mark 5C) and cannot support 16 Gbps and higher data rates without many systems operating in parallel, which is both expensive and cumbersome. The Mark 6 system, based fully on commercial-off-the-shelf (COTS) hardware, is being jointly developed by MIT Haystack Observatory in collaboration with Conduant Corporation and the NASA/GSFC High-End Network Computing Group.

2. Mark 6 Goals

The goals of the Mark 6 system include:

- 16 Gbps sustained record and playback capability
- ≥ 32 Gbps burst-mode capability
- General Ethernet packet recorder (can be straight-forwardly adapted to other interfaces as well)
- Use of inexpensive high-performance commercial-off-the-shelf (COTS) hardware as a basis
- Ability to be easily upgraded on Moore's Law curve

- Linux OS with open-source software
- Playback as standard Linux files
- Resilient File System to manage slow and failed disks
- e-VLBI support
- Smooth transition from Mark 5
- Preservation of investment in existing Mark 5 systems and disk libraries as much as possible

3. Main Characteristics of the Mark 6 System

3.1. Input Data Interfaces

The Mark 6 can accommodate data from up to four 10GigE sources, operating at a maximum aggregate of 16 Gbps. Each input may operate at a different data rate. Normally, data are transmitted to the Mark 6 in a UDP packet stream to minimize load on the Mark 6 Ethernet NIC cards and sustain the highest possible data rate.

3.2. Hardware

In contrast to the Mark 5 system, the Mark 6 controller incorporates only COTS hardware, although components are carefully selected for performance and compatibility with other system elements. A high-end motherboard, CPU, and RAM memory are used to maximize performance; a standard Mark 6 system is outfitted with ~12 GB of high-speed RAM memory, although up to ~128 GB can be accommodated for burst-mode applications. Outwardly, a Mark 6 system looks much like a Mark 5 system, but with the following differences:

- Four Mark 6 disk modules are required to support 16 Gbps.
- Mark 6 disk modules look similar to Mark 5 disk modules except that the data connections are via front-panel cables.

A Mark 6 prototype system is shown in Figure 1. The data cables emerge from a 1U cable-management panel between the Mark 6 “system chassis” (all data electronics) and Mark 6 “expansion chassis” (contains only a power supply for the associated disk modules). The cables can easily be pushed back into the cable-management panel when not being used.

3.3. Mark 6 Disk Modules

The Mark 6 system supports only SATA-interface data disks. Mark 6 disk modules are similar to Mark 5 disk modules (eight disks per module) except that each Mark 6 disk module connects to the disk controller via two COTS SAS cables. Each SAS cable supports four disks, so that a module requires the connection of two such cables. Two cables are assigned to each disk-module slot, but they may be connected to the associated module in any order. Power to the module is provided through a backplane connector on the module.

The use of external-SATA cables to connect the Mark 6 controller to the modules is different from Mark 5, but it is quite workable and actually has some advantages. The heaviest wear is on the cable ends that connect to the modules; when a cable connector wears out, the cable is easily and inexpensively replaced.



Figure 1. Photograph of Mark 6 development system.

A set of eight LEDs alongside the data connectors gives a visual indication of disk activity. An inexpensive kit is available to convert existing Mark 5 SATA modules to Mark 6 modules; the conversion kit consists of a new module backplane and front-panel. Converted modules are not backwards compatible with Mark 5 systems.

3.4. Software

The Mark 6 system operates under a standard Linux OS. Application software is written primarily in C, C++, and Python. The Mark 6 performance arises from a technique that decouples network and disk datastreams with elastic buffers and kernel-bypass DMA for data transfer.

3.5. Monitor & Control and Disk Management

A simple VSI-S command set has been specified for the Mark 6.

Depending on the recording data rate, different numbers of simultaneously-operating disk modules are required. A single 8-disk module with modern disks will support 4 Gbps; two modules (16 disks) are required for 8 Gbps, four modules (32 disks) are required for 16 Gbps. In order to accommodate these different requirements, the concept of a ‘group’ has been created to identify the collection of one or more disk modules needed to support a particular observing requirement. A ‘group’ is created by ‘bonding’ a specified set of modules together for the duration of a particular data set. When the individual modules of a group are re-initialized, they are released from the group and become individual again.

3.6. Data Format on Disks and Playback

The Mark 6 writes standard Linux files to the set of data disks. In order to maintain the maximum recording-rate capability, the Mark 6 records the entire Ethernet packet from all data streams and is oblivious to the format of the actual data frame within the Ethernet packet. On replay, Ethernet header and trailer data are normally stripped so that the user sees only the actual payload packet (VDIF or Mark 5B, etc).

3.7. e-VLBI Capabilities

Full e-VLBI capabilities are planned for the Mark 6.

3.7.1. 16 Gbps Mark 6 VLBI Demonstration

On 24 October 2011, a 16-Gbps VLBI test experiment using a pair of development Mark 6 units was conducted between the Westford 18-m antenna and the 12-m VLBI2010 development antenna at NASA/GSFC in Maryland. As shown in Figure 2, a single 500 MHz-wide IF signal was duplicated (in analog) eight times, then processed through RDBE digital backend subsystems to produce an aggregate 16 Gbps (on four 10GigE data lines) from each station.

The data were cross-correlated by the Haystack DiFX software correlator to produce normal fringes on all eight 500MHz-wide channels. The fringe-fitting results from one of the eight channels is shown in Figure 3; the results from the other channels are essentially identical.

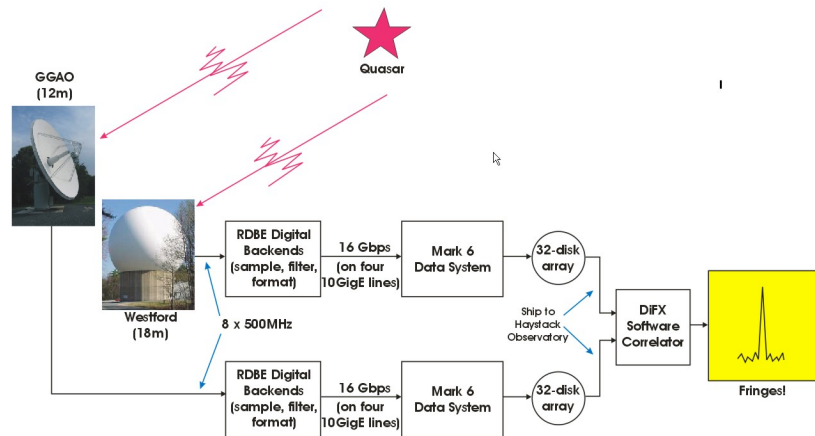


Figure 2. Block diagram of 16-Gbps VLBI demonstration experiment between the Westford (MA) and GGAO (MD) antennas. Due to receiver limitations, a 500 MHz BW IF was duplicated eight times to provide the analog data required to create a 16 Gbps data stream at each antenna.

4. Upgrading from Mark 5 to Mark 6

A Mark 6 system is made up of two chassis—a ‘system chassis’ (which contains all of the data electronics) and an ‘expansion chassis’ (which contains only a power supply)—each chassis capable of mounting two Mark 6 disk modules. The bare chassis for both the ‘system’ and the ‘expansion’ chassis are identical to a Mark 5 chassis. Both chassis must be upgraded with new chassis backplanes with power connectors to mate to the Mark 6 modules. By upgrading with the new chassis backplanes and adding the appropriate data electronics, a Mark 5 chassis can be converted to a Mark 6 “system chassis”; the “expansion chassis” requires only an appropriate power supply. In addition, an inexpensive ‘cable management tray’ must be procured for data-cable management. Upgrade kits to convert Mark 5 systems to Mark 6, as well as Mark 5 SATA disk-module conversion kits, are available from Conduant Corporation, as are complete new systems.




Figure 3. Correlation output results for one of the eight 500 MHz channels; note that the cross-power spectrum shows only ~ 400 MHz of actual noise bandwidth was available, consistent with the actual signal bandwidth produced by the receiver.

5. Summary

The Mark 6 VLBI data system is a major step forward in data-rate capability over previous VLBI data system, and is the first high-performance system to use fully COTS data hardware. The high performance and low cost of the Mark 6 system, as well as future improvements due to the normal progress of COTS technology, will help to maintain its long-term viability. In the short term, geodetic-VLBI and mm-VLBI will be the major beneficiaries of this new capability, but in the longer term it will also enable much higher sensitivity over a broad range of VLBI applications. The Mark 6 system is continuing to be developed and is expected to be available to the general VLBI community in late 2012. More information about the Mark 6 system is available at <http://www.haystack.edu/tech/vlbi/mark6/index.html>.